Improving Wireless Simulation through Noise Modeling

HyungJune Lee and Philip Levis
Computer Systems Laboratory
Stanford University

Alberto Cerpa
School of Engineering
UC Merced

Never What You Hoped

- Simulation never seems to capture reality
- Design a protocol, simulate it heavily, deploy it...
  ... and it doesn't work.
- Decades of work on analytical models from the RF community
- Recent work on empirical models from the systems community
- What's missing?

Network Models

- Link layer down
  - Rates from distances [TOSSIM]
  - Rates from network [EmStar]

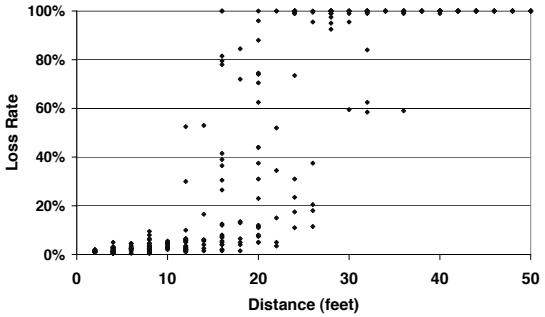

- Physical layer up
  - Bit error rate from signal-to-noise ratio
  - Signal is complicated
    - Varies with TX signal strength [Son et al.]
    - Signal is asymmetric [Zuniga et al.]
    - Noise floors vary [Srinivasan et al.]

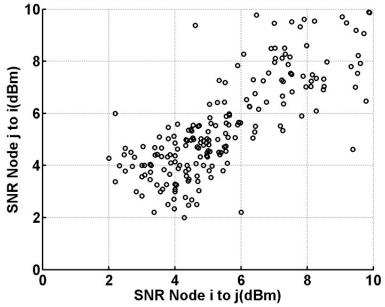

Losses are not Independent

- Simulators model independent packet losses
- Real networks observe correlated losses
- Causes?
  - Long-term variations from signal changes
  - Short-term variations from external noise

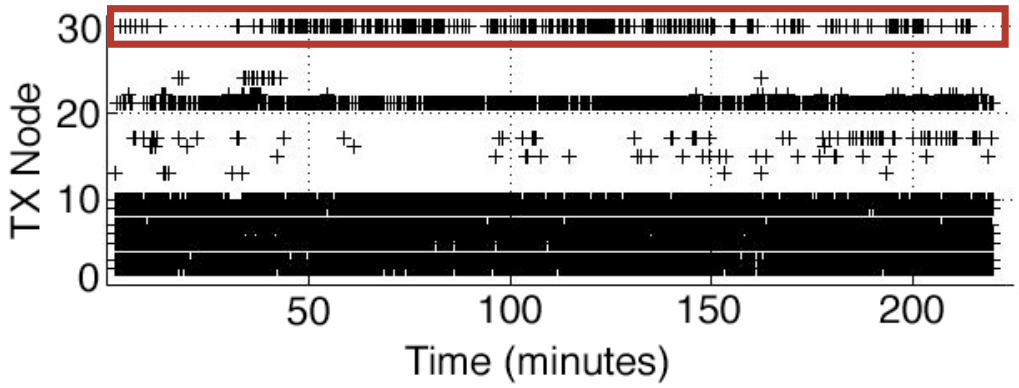

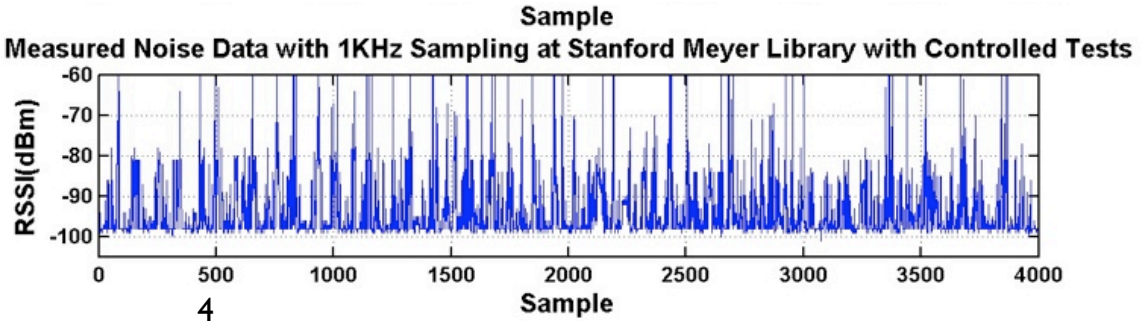

Summary of Contribution

- Capture high-frequency noise traces from a real network.
- Explore three noise models.
- CPM reduces KW distance between delivery distributions from 0.3 to 0.04.
- Improvement affects network protocol behavior.
- Modeling external noise captures important real-world behavior.

Outline

- Improving simulation
- Noise sampling and simulation
- Results
- Next steps

Outline

- Improving simulation
- Noise sampling and simulation
- Results
- Next steps

Signal to Noise

- Experimentally derive a CC2420 SNR curve
- Shielded micaZ nodes with variable attenuator

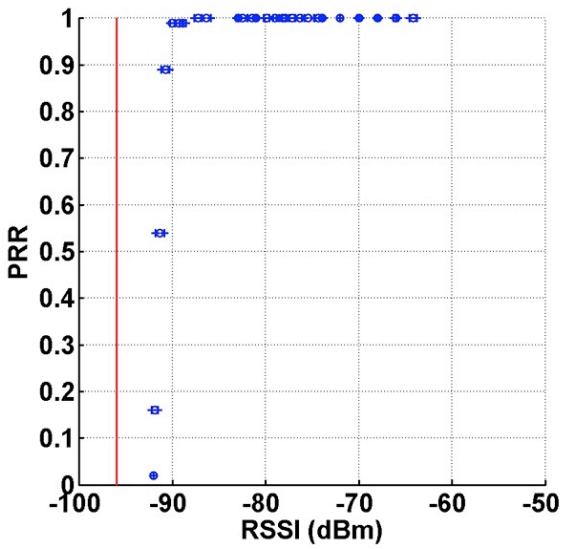

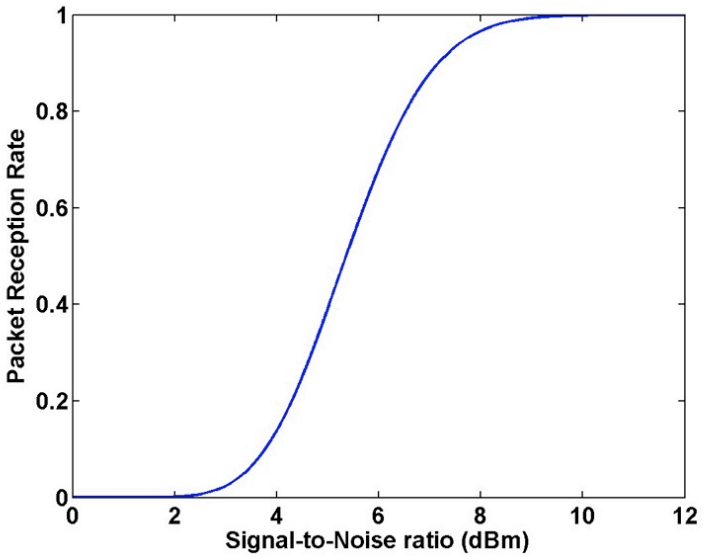

Noise Sampling

- Sample RF energy at a frequency greater than maximum packet speed
- Implemented on CC2420 radio by reading RSSI register at 1kHz
- Store trace to flash (~3 min, 1.95×10^5 samples)
- Output to serial port when done

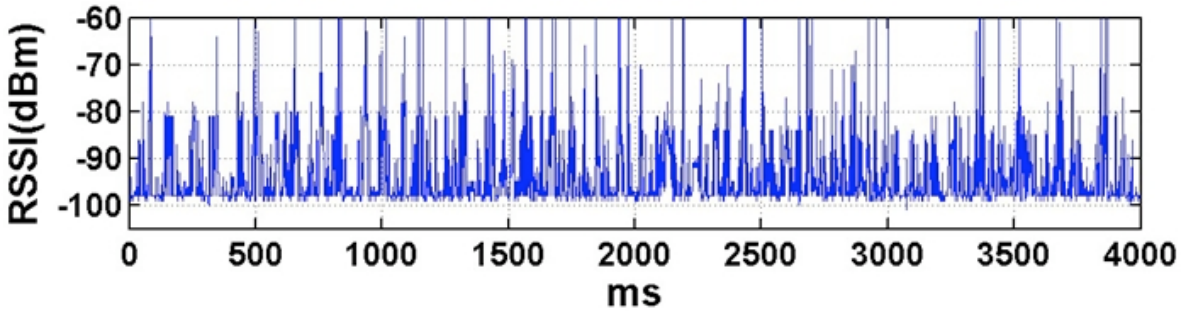

Meyer Library,
next to active laptop

What Links Look Like

- Need to evaluate whether simulation is representative of reality
- Metric: conditional packet delivery function (CPDF)
- Example: intermediate link (PRR=51%)
  - IPI=15ms, heavy Meyer trace

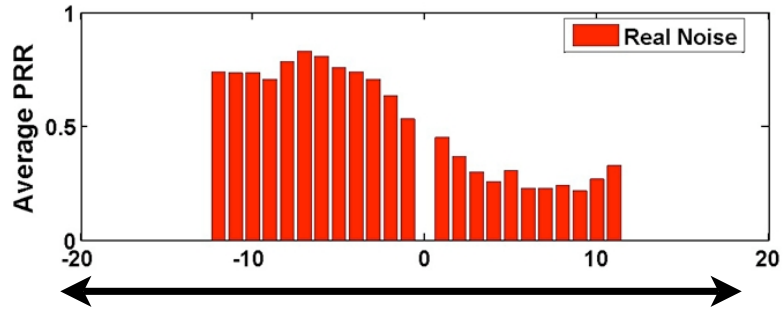

Successes Failures

Good, Bad, and Ugly

- Adjust RSSI to achieve PRR, replay noise trace

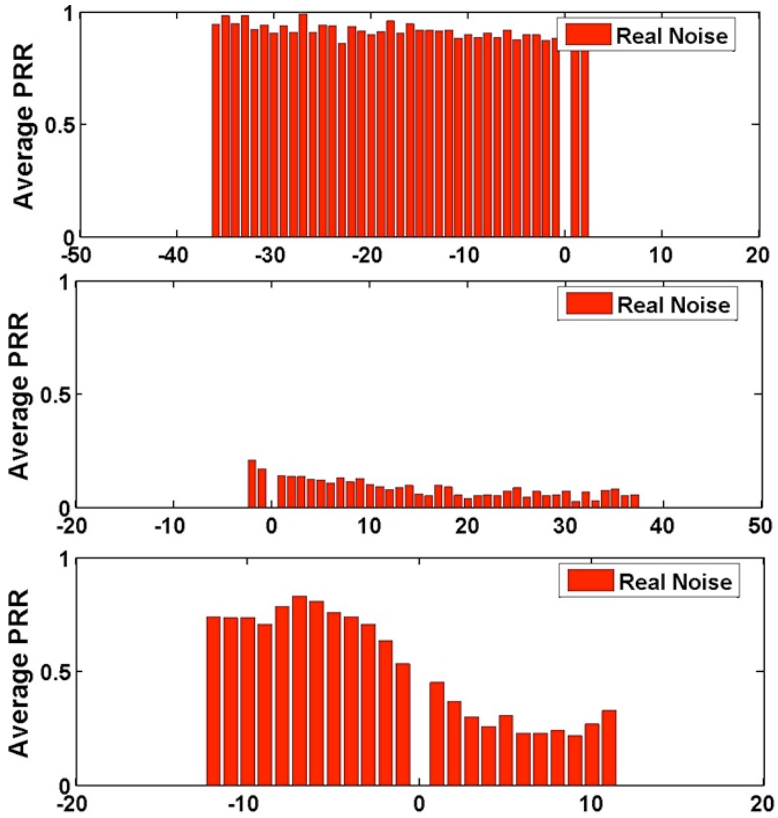

Good (PRR=90%)

Bad (PRR=11%)

Ugly (PRR=51%)

Three Algorithms

- Naive
- Correlation distortion method
- Closest pattern matching (CPM)

Three Algorithms

- Naive
- Correlation distortion method
- Closest pattern matching (CPM)

CPM Initialization

- Scan noise trace, keeping a history of size *k*
- For each signature of *k* prior noise readings, construct probability distribution of next reading

0	2	1	2	0	2	2	0	0	1	1	1	1	2	0	0	0	2	9	0

signature	0	1	2	9
00	33%	33%	33%	0%
01	0%	100%	0%	0%
02	0%	33%	33%	33%
11	0%	66%	33%	0%
12	100%	0%	0%	0%
20	66%	0%	33%	0%
21	0%	0%	100%	0%
22	100%	0%	0%	0%

$k = 2$

CPM Initialization

- Scan noise trace, keeping a history of size *k*
- For each signature of *k* prior noise readings, construct probability distribution of next reading

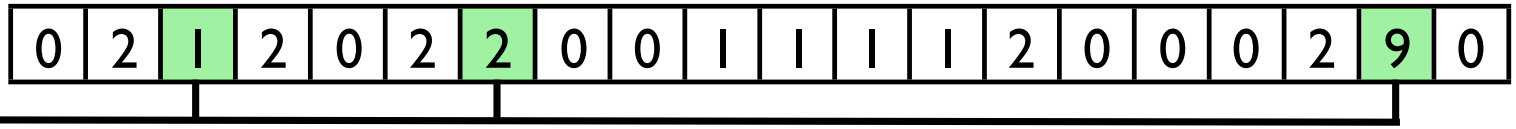

signature	0	1	2	9
00	33%	33%	33%	0%
01	0%	100%	0%	0%
02	0%	33%	33%	33%
11	0%	66%	33%	0%
12	100%	0%	0%	0%
20	66%	0%	33%	0%
21	0%	0%	100%	0%
22	100%	0%	0%	0%

$k = 2$

Running CPM

- Take first *k* noise values directly from trace
- For each subsequent value, sample from distribution matching the prior *k* values
- If there is no match for the prior *k*, sample from most common distribution
- Edge cases
  - *k = 0* samples independently
  - *k* = |*T*| replays trace T exactly

Outline

- Improving simulation
- Noise sampling and simulation
- Results
- Next steps

Meyer Heavy Trace

Real Noise

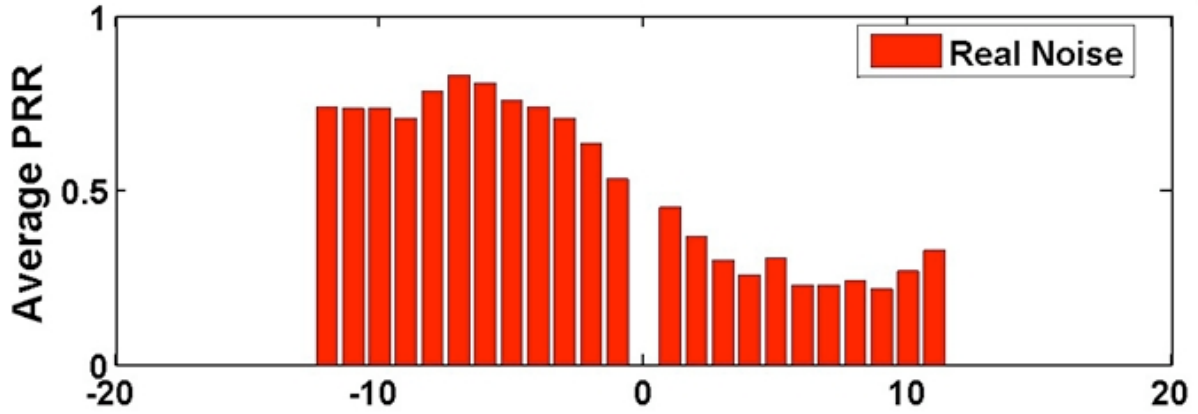

Naive Sampling

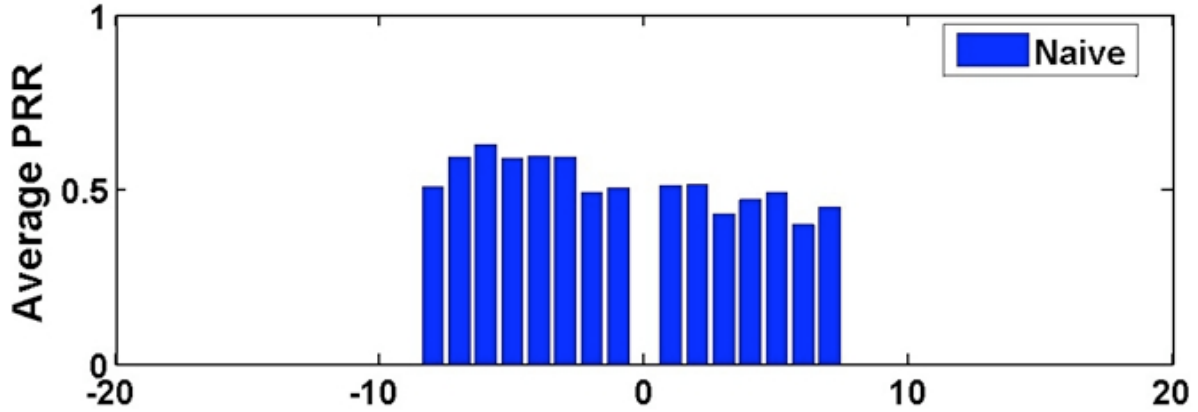

Successes ← → Failures

Meyer Heavy Trace

Real Noise

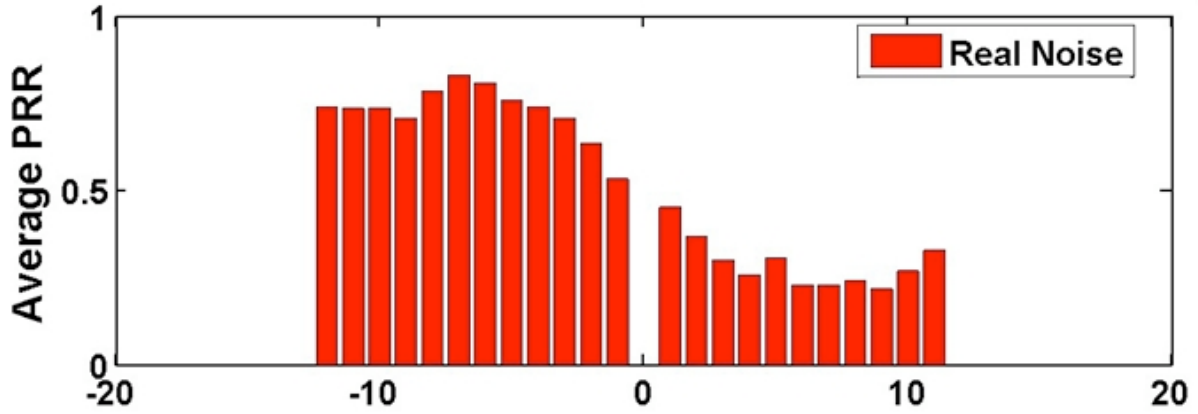

TOSSIM TinyOS 1.x

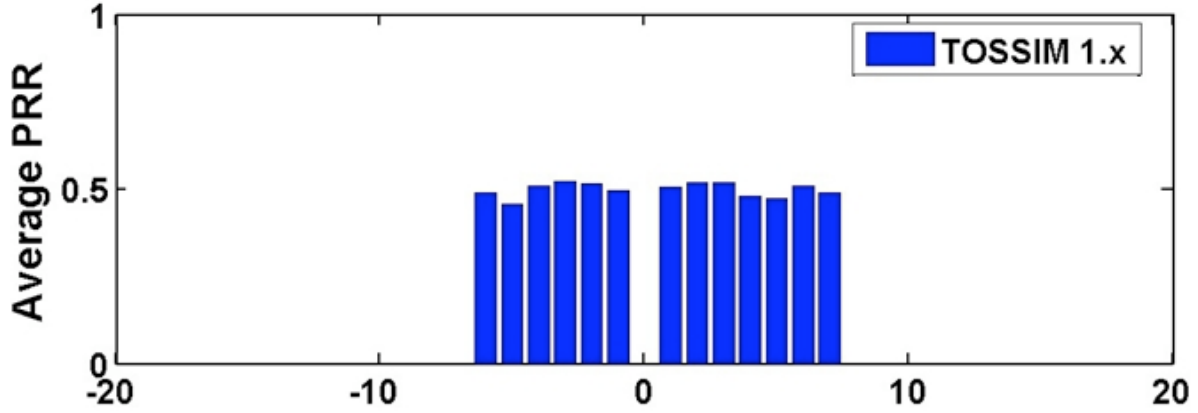

Successes ⟷ Failures

Meyer Heavy Trace

Real Noise

TOSSIM TinyOS 2.0

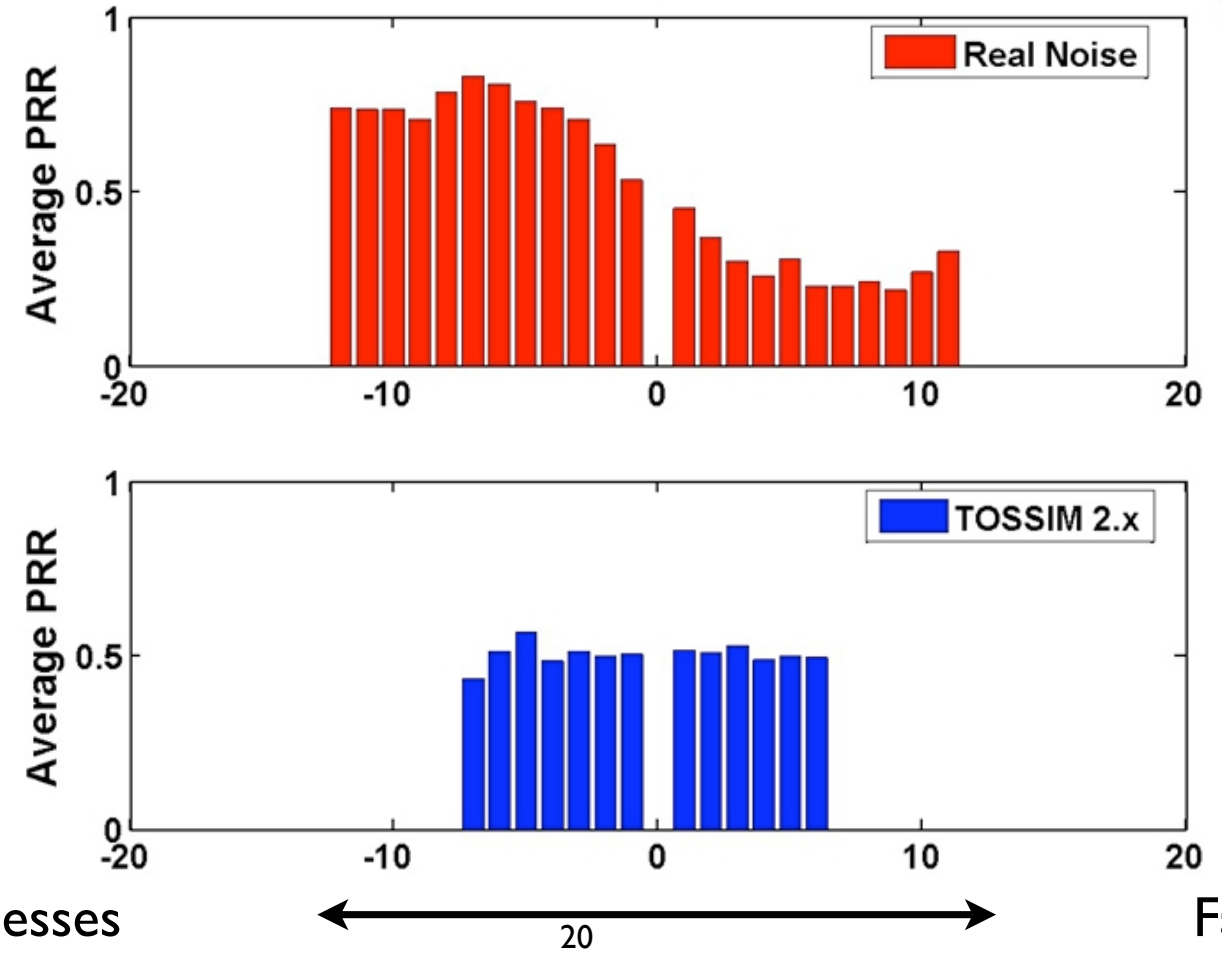

Successes

20

Failures

Meyer Heavy Trace

Real Noise

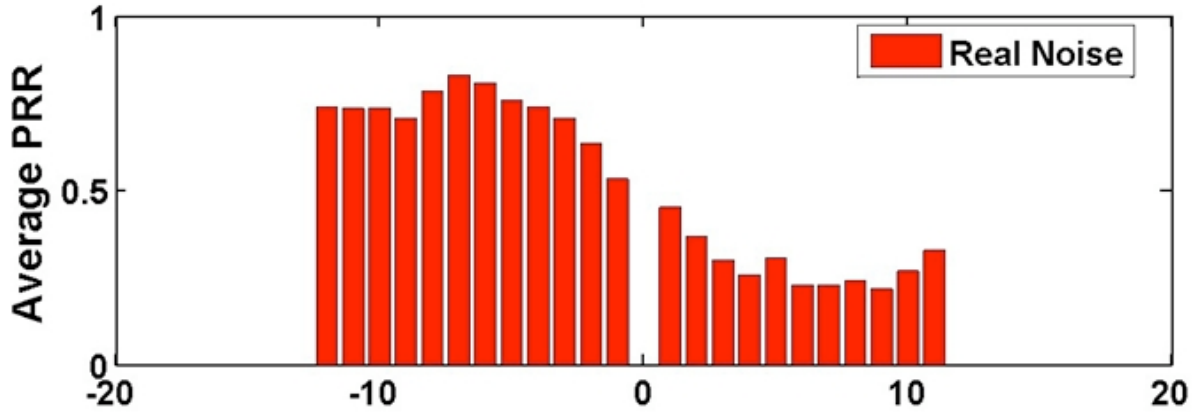

EmStar

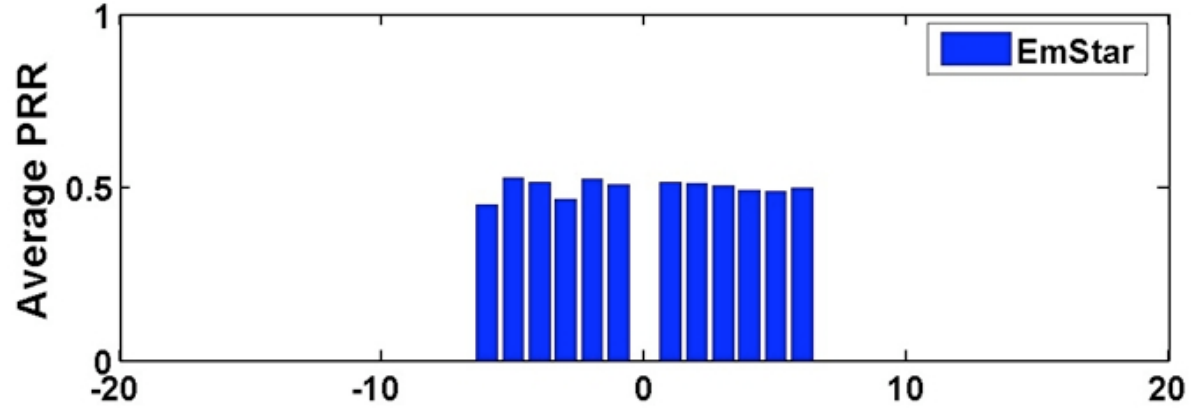

Successes ⟷ Failures

Meyer Heavy Trace

Real Noise

ns2

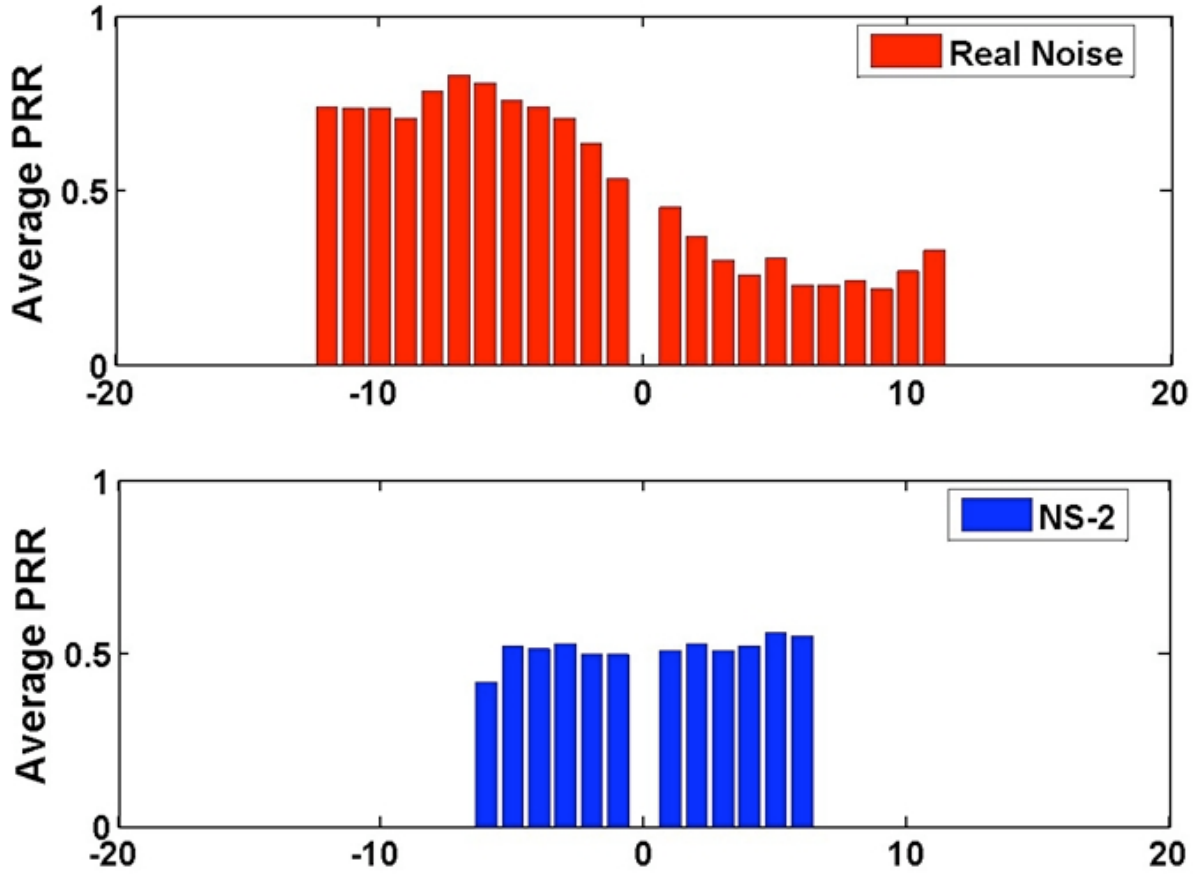

Successes

Failures

Meyer Heavy Trace

Real Noise

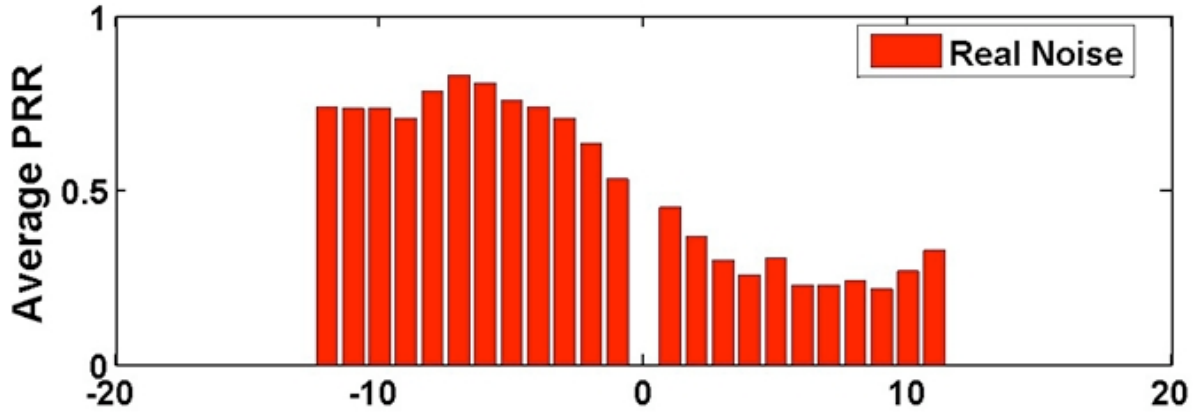

CPM

k=20
Q=5

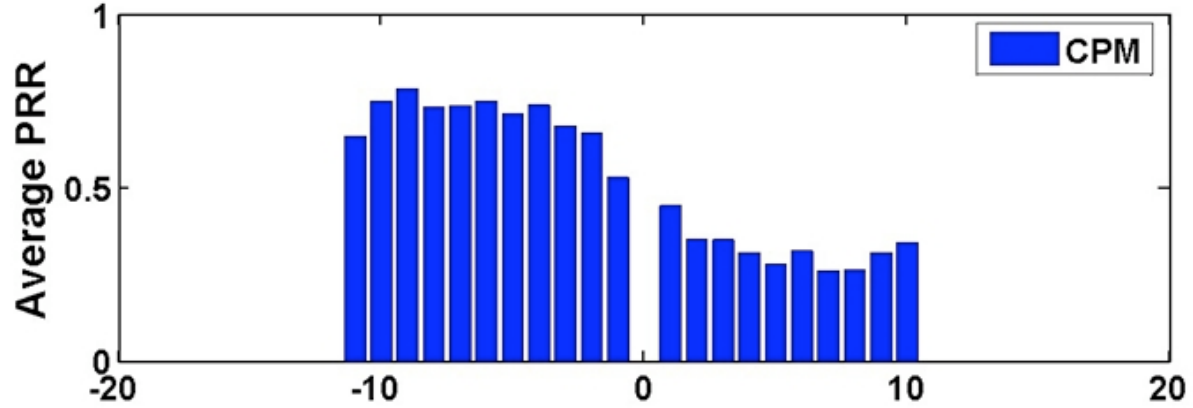

Successes ⟷ Failures

Protocol Effects

- Link estimator in TinyOS 2.0 collection tree protocol uses acknowledgments to measure ETX
- Bursts of losses can cause rapid link value changes
- Measure how many times protocol switches parents
- Using CPM doubles changes over other methods

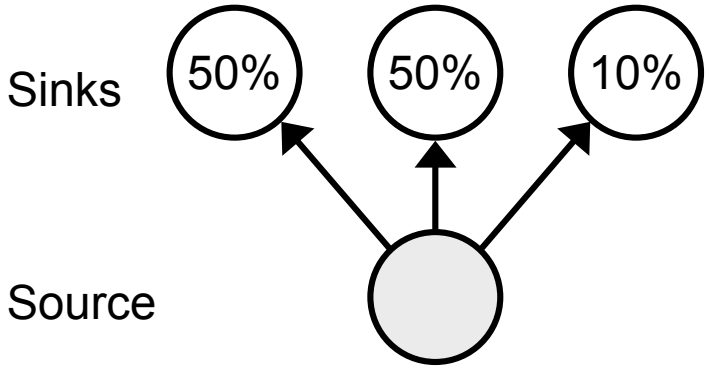

Outline

- Improving simulation
- Noise sampling and simulation
- Results
- Next steps

CPM

- By considering statistical correlation between noise values, CPM greatly improves simulation quality
- Can capture bursts of losses as observed in real networks, as well as other short-term link variations
- Reduces KW distance from real noise to within 0.04
- These improvements affect network protocol behavior

Future Work

- CPM only considers a node's noise in isolation
  - Real nodes share observations
  - Next step: spatial correlation
- Long term link variations
  - Experimental results indicate this is due to RSSI shifts
  - Next step: model RSSI trends

Current State

- CPM allows us to simulate the behavior of real networks
- Incorporated into TOSSIM in TinyOS 2.0.1
  - Being released this weekend, at the TTX
- We have written a TinyOS application that samples 1kHz noise traces
  - `tinyos-2.x-contrib/stanford-sing/apps/RssiSample`
- If you sample your environments and send us your traces, so we can include them in TOSSIM

Questions

Running CPM Example

0	2				

signature	0	1	2	9
00	33%	33%	33%	0%
01	0%	100%	0%	0%
02	0%	33%	33%	33%
11	0%	66%	33%	0%
12	100%	0%	0%	0%
20	66%	0%	33%	0%
21	0%	0%	100%	0%
22	100%	0%	0%	0%

$k = 2$

Running CPM Example

0	2	2			

signature	0	1	2	9
00	33%	33%	33%	0%
01	0%	100%	0%	0%
02	0%	33%	33%	33%
11	0%	66%	33%	0%
12	100%	0%	0%	0%
20	66%	0%	33%	0%
21	0%	0%	100%	0%
22	100%	0%	0%	0%

$k = 2$

Running CPM Example

0	2	2	0		

signature	0	1	2	9
00	33%	33%	33%	0%
01	0%	100%	0%	0%
02	0%	33%	33%	33%
11	0%	66%	33%	0%
12	100%	0%	0%	0%
20	66%	0%	33%	0%
21	0%	0%	100%	0%
22	100%	0%	0%	0%

$k = 2$

Running CPM Example

0	2	2	0	0	

signature	0	1	2	9
00	33%	33%	33%	0%
01	0%	100%	0%	0%
02	0%	33%	33%	33%
11	0%	66%	33%	0%
12	100%	0%	0%	0%
20	66%	0%	33%	0%
21	0%	0%	100%	0%
22	100%	0%	0%	0%

$k = 2$

Running CPM Example

0	2	2	0	0	1

signature	0	1	2	9
00	33%	33%	33%	0%
01	0%	100%	0%	0%
02	0%	33%	33%	33%
11	0%	66%	33%	0%
12	100%	0%	0%	0%
20	66%	0%	33%	0%
21	0%	0%	100%	0%
22	100%	0%	0%	0%

$k = 2$

Running CPM Example

0	2	2	0	0	1

0	2	1	2	0	2	2	0	0	1	1	1	1	2	0	0	0	2	9	0

signature	0	1	2	9
00	33%	33%	33%	0%
01	0%	100%	0%	0%
02	0%	33%	33%	33%
11	0%	66%	33%	0%
12	100%	0%	0%	0%
20	66%	0%	33%	0%
21	0%	0%	100%	0%
22	100%	0%	0%	0%

$k = 2$

Quantization

- Given a range of R possible RSSI readings, state space of prior readings is R^k
- Very sparse state spaces cause distributions to have few values, leading to low match rates
- Quantize noise readings by Q:

$$q_i = \left\lfloor \frac{n_i}{Q} \right\rfloor$$

- Choose value in range $Q \cdot [q_i, q_i+1]$